arXiv:1702.03967v1 [math.DS] 13 Feb 2017

Nonlinear Kalman Filtering for Censored Observations

Joseph Arthur[a], Adam Attarian[b], Franz Hamilton[c,*], Hien Tran[c]

[a] *Department of Statistics, Stanford University, Stanford, California, 94305 USA*
[b] *MIT Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Massachusetts, 02420 USA*
[c] *Department of Mathematics and Center for Quantitative Sciences in Biomedicine, North Carolina State University, Raleigh, North Carolina, 27695 USA*

Abstract

The use of Kalman filtering, as well as its nonlinear extensions, for the estimation of system variables and parameters has played a pivotal role in many fields of scientific inquiry where observations of the system are restricted to a subset of variables. However in the case of censored observations, where measurements of the system beyond a certain detection point are impossible, the estimation problem is complicated. Without appropriate consideration, censored observations can lead to inaccurate estimates. Motivated by the work of [1], we develop a modified version of the extended Kalman filter to handle the case of censored observations in nonlinear systems. We validate this methodology in a simple oscillator system first, showing its ability to accurately reconstruct state variables and track system parameters when observations are censored. Finally, we utilize the nonlinear censored filter to analyze censored datasets from patients with hepatitis C and human immunodeficiency virus.

Keywords: extended Kalman filter, censored observation, parameter estimation, hepatitis C virus (HCV), human immunodeficiency virus (HIV)

1. Introduction

The use of data assimilation for the estimation of unobserved model variables and parameters has become standard practice in modern scientific analysis.

*Corresponding author
Email address: fwhamilt@ncsu.edu (Franz Hamilton)

Kalman filtering [2] and its nonlinear extensions such as the ensemble Kalman filter and extended Kalman filter have gained increasing popularity in application to a variety of problems arising from the physical and biological sciences [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16].

Of recent interest in the field of biomedicine has been the use of ordinary differential equations to model viral infection dynamics such as human immunodeficiency virus (HIV) and hepatitis C virus (HCV) [17, 18, 19]. Such models can provide insights into disease behavior, treatment, and ultimately improve patient outcomes. Their use for the development of patient specific treatment regimens remains an exciting possibility. However, these models are parameterized by a number of unknown parameters and observation of the system is limited to a noisy subset of the dynamical variables.

Several methodologies have developed to handle this problem of state and parameter estimation from noisy observations. In particular, the use of Kalman filtering for joint state and parameter estimation has been the topic of several recent papers [20, 21, 22, 23]. Unfortunately, this estimation process is further complicated when we consider that the assays used in viral studies for data collection often have a detection limit beyond which accurate measurements are impossible. We refer to these data as *censored*. Measurements within the detection limit are considered *uncensored*. Ignoring the censored data can lead to bias in the estimates [24]. As such, a proper framework for handling censored observations is required.

Kalman filtering for censored observations has been the topic of several recent works [1, 24, 25]. Of particular interest is the method proposed in [1]. There the authors derived an auxiliary set of equations for the Kalman filter which provided a modified Kalman gain and covariance update formula to allow for correct inference given censored measurements. The underlying assumption though was that the system of interest is linear. Unfortunately the majority of physical systems and the models representing them are nonlinear, such as those describing the dynamics of HCV and HIV. Our goal in this article is to extend the methodology presented in [1] to the case of nonlinear system dynamics. We

derive a modified version of the extended Kalman filter allowing for the accurate joint estimation of state variables and parameters in nonlinear systems in the presence of censored data.

We validate our proposed nonlinear censored filter first in a synthetic oscillator system where a detection limit for system observation is imposed. We show the fidelity of filter's state variable and parameter reconstruction even when we have partial observability of the system and several of the data are censored. Additionally, we demonstrate the capability of the filter to track system non-stationarity in the form of a drifting parameter whose dynamics are unknown. Motivated by our success in this synthetic example, we consider the difficult problem of state and parameter estimation for clinical viral data. In particular we examine two datasets from an HCV and HIV clinical study, both of which contain numerous censored data in their respective viral load measurements.

In analyzing these clinical datasets, we follow very closely the work done in [26] and [27] for the HCV and HIV data repsectively. There, the authors provided a detailed model identifiability analysis for these datasets and performed estimation using the expectation maximization algorithm [28]. Our belief is that the filter should not provide more reliable or accurate estimates than those calculated by expectation maximization, in fact they should be comparable. Therefore we treat the results of [26, 27] as "ground truth" and aim to show that the proposed nonlinear censored filter is able to reproduce similar estimates. The true utility of the filter is that it provides sequential estimation allowing for the online joint estimation of state variables and parameters and the possibility of tracking parameters whose values drift over time, both of which expectation maximization are unable to do. These capabilities are of particular interest in the field of personalized medicine where researchers may be analyzing clinical data whose measurements span over several months or years and an accurate and timely estimate of the current system state is necessary for appropriate treatment or intervention.

2. Nonlinear Kalman Filtering with Censored Observations

We assume the following nonlinear system with continuous-time state dynamics and discrete observations

$$\begin{aligned} \dot{x}(t) &= f(t,x) + w(t) \\ z(t_k) &= h(x(t_k)) + v_k, \end{aligned}$$

where x is an n dimensional state vector and z is an m dimensional observation vector. w and v are Gaussian noise terms with covariances Q and R respectively. The estimation of Q and R is key to the success of any filtering methodology. Here, we perform offline tuning of these error covariance matrices to obtain optimal filter performance.

Due to the system nonlinearity, the standard Kalman filter can not be applied directly. Several nonlinear filters have developed, such as the ensemble Kalman filter (EnKF) and extended Kalman filter (EKF) [29, 30]. Here we focus solely on the EKF, which performs a linearization of the system dynamics at each step of the filter. For a detailed derivation of the algorithm see [31].

The EKF is a sequential estimator that consists of a prediction and update step. We solve the following system

$$\begin{aligned} \dot{\hat{x}} &= f(t,\hat{x}) \\ \dot{P} &= PF^T + FP + Q, \end{aligned}$$

with initial conditions $\hat{x}_{k-1}$ and P_{k-1} from t_{k-1} to t_k to compute $\hat{x}_k^-$ and P_k^-, our prior state and covariance matrix estimates. F is the linearization of the system dynamics, namely $F = \nabla f(\hat{x})$. We form the linearization of the observation operator, $H_k = \nabla h(\hat{x}_k^-)$, and then implement the standard Kalman update equations to correct our state and covariance estimates

$$\begin{aligned} \hat{x}_k =& \hat{x}_k^- + K_k \left[z_k - h(\hat{x}_k^-) \right] \\ P_k =& \left[I - K_k H_k \right] P_k^- \\ K_k =& P_k^- H_k^T \left[H_k P_k^- H_k^T + R \right]^{-1}. \end{aligned}$$

2.1. Filtering with Censored Data

In the case of censored data, where the true value of the observation beyond a certain lower or upper detection limit is unknown, the estimation problem is complicated. Treating these censored observations as uncensored measurements leads to inaccurate estimates. In [1], Gabardós and Zufiria addressed this problem of state estimation in the presence of censored data in the Kalman filter framework. The authors derived a new set of equations for the filter which appropriately accounts for censored data during the Kalman update step. In this article we extend these ideas to the nonlinear case, deriving an auxiliary set of update equations for the EKF to accurately handle censored data. The derivation included here follows very closely that in [1], though our assumption throughout is that our system of interest is nonlinear.

We use U_k to denote the vector of all uncensored observations up to time t_k. Similarly, let C_k denote the vector of censored observations, each of which lies in some possibly infinite interval $\mathcal{Z}$. For simplicity, we will write $C_k \in \mathcal{Z}$. The filter proceeds at every step by first estimating the state and error covariance ignoring any censored observations. We denote these naive estimates with $\hat{x}_{k(uc)}$ and $P_{k(uc)}$ and use $\hat{x}_k$ and P_k to denote the final estimates, which are additionally conditioned on the censored observations lying in $\mathcal{Z}$. To calculate the naive estimates, we use a modified gain term:

$$K_k = \begin{cases} 0 & \text{if } z_k \text{ is censored} \\ P^-_{k(uc)} H_k^T \left[H_k P^-_{k(uc)} H_k^T + R \right]^{-1} & \text{otherwise.} \end{cases}$$

Therefore when z_k is censored, we have $\hat{x}^-_{k(uc)} = \hat{x}_{k(uc)}$ and $P^-_{k(uc)} = P_{k(uc)}$, i.e., the predicted values are equal to the naive estimates.

In the case of a censored observation, we calculate the mean and approximate error covariance for the censored observation conditional on the uncensored data, namely

$$\begin{aligned} \hat{C}_{k(uc)} =& h(\hat{x}_{k(uc)}) \\ P^C_{k(uc)} =& \left[H_k P_{k(uc)} H_k^T + R \right], \end{aligned}$$

We also compute

$$P^{Cx}_{k(uc)} = H_k P_{k(uc)},$$

the covariance between the censored observation and the state. Using multivariate Gaussian calculations (see Appendix), the final state and covariance update equations are defined as

$$\hat{x}_k = \hat{x}_{k(uc)} + K'_k \left[\hat{C}_k - \hat{C}_{k(uc)}\right] \tag{1}$$

$$P_k = P_{k(uc)} - K'_k \left[P^C_{k(uc)} - P^C_k\right] (K'_k)^T, \tag{2}$$

where the new gain term is

$$K'_k = P^{xC}_{k(uc)} \left(P^C_{k(uc)}\right)^{-1} \tag{3}$$

and

$$P^{xC}_{k(uc)} = \left(P^{Cx}_{k(uc)}\right)^T.$$

Note that $\hat{C}_k$ and P^C_k are the mean and covariance of the censored observation given the uncensored observations and conditioned on the censored observation lying in $\mathcal{Z}$. This computation is done using the `tmvtnorm` package in R, which computes the mean and covariance of truncated multivariate normal random variables [32]. After the first censored observation, (1), (2), and (3) are used as the state and covariance update equations. Additionally though, we must update $\hat{C}_{k(uc)}, P^C_{k(uc)}$, and $P^{Cx}_{k(uc)}$ at every step of the filter. This update is carried out in two ways, depending on whether or not z_k is censored.

In the censored case we first update the covariance $P^{Cx}_{k-1(uc)}$ to account for the change in state from t_{k-1} to t_k. Momentarily abbreviating this covariance as D, we solve the system

$$\dot{D} = DF^T \tag{4}$$

$$\dot{\hat{x}} = f(\hat{x}) \tag{5}$$

from t_{k-1} to t_k with initial conditions $D(t_{k-1}) = P^{Cx}_{k-1(uc)}$ and $\hat{x}(t_{k-1}) = \hat{x}_{k-1}$. The result of this computation is that $D(t_k)$ is approximately the covariance between C_{k-1} and x_k, conditional on only the uncensored observations (see

Appendix for details). We call this covariance $P^{Cx}_{k-1,k(uc)}$ and compute the final updated covariance as

$$P^{Cx}_{k(uc)} = \begin{bmatrix} P^{Cx}_{k-1,k(uc)} \\ H_k P_{k(uc)} \end{bmatrix}.$$

Now, we update the naive covariance of the censored observations as

$$P^{C}_{k(uc)} = \begin{bmatrix} P^{C}_{k-1(uc)} & P^{Cx}_{k-1,k(uc)} H_k^T \\ H_k (P^{Cx}_{k-1,k(uc)})^T & P^{z}_{k(uc)} \end{bmatrix},$$

where the covariance of the new observation is

$$P^{z}_{k(uc)} = \left[H_k P_{k(uc)} H_k^T + R \right].$$

Similarly, updating the naive estimate for the censored observations gives

$$\hat{C}_{k(uc)} = \begin{bmatrix} \hat{C}_{k-1(uc)} \\ h(\hat{x}_{k(uc)}) \end{bmatrix}.$$

In the case that z_k is not censored, the calculations become slightly more complicated. We first use equations (4) and (5) to compute $P^{Cx-}_{k-1,k(uc)}$, which is equivalent to $P^{Cx-}_{k(uc)}$ since $C_k = C_{k-1}$. This predictive covariance can be updated as

$$P^{Cx}_{k(uc)} = P^{Cx-}_{k(uc)} \left[I - H_k^T K_k^T \right].$$

The naive expectation of the censored data vector can be updated according to

$$\hat{C}_{k(uc)} = \hat{C}_{k-1(uc)} + P^{Cx-}_{k(uc)} H_k^T (P^{z}_{k(uc)})^{-1} \left[z_k - h(\hat{x}^{-}_{k(uc)}) \right],$$

which is analogous to the state update equation in the basic Kalman filter. Similarly, we use the equation

$$P^{C}_{k(uc)} = P^{C}_{k-1(uc)} - P^{Cx-}_{k(uc)} H_k^T (P^{z}_{k(uc)})^{-1} H_k (P^{Cx-}_{k(uc)})^T$$

to update the error covariance for the censored observations.

Of course, with an increasing number of censored data the above algorithm can become computationally unwieldy due to the increasing dimension of the covariance matrices. In [1] the authors reason that previous censored data can

be forgotten over time, allowing for a reduction in the algorithm's computational complexity. In particular the columns of the modified gain term K'_k defined in (3), where each column corresponds to a censored observation, will naturally decay over time to 0 as more data is processed. Additionally if there are a sufficient number of uncensored observations after a censored measurement, the correlation between the censored observation and the state becomes very small. With these ideas in mind, we can introduce approximations to the state and covariance update by removing past censored observations. This in effect reduces the computational complexity of the algorithm and would allow us to only use subsets of the censored observations for a period of time.

3. State and Parameter Estimation in Oscillator System

As an demonstrative example, we consider the estimation problem in the following oscillator system

$$\begin{aligned} \dot{x}_1 &= \alpha x_2 \\ \dot{x}_2 &= 4 - 4x_1, \end{aligned}$$

where α is an arbitrary parameter. Our assumption is that only noisy observations of x_1 sampled at rate $dt = 0.2$ are available. Observations of x_1 though are restricted in that any measurement below a value of 0.8 is censored, implying a censored interval of $[-\infty, 0.8]$. Given these noisy censored observations, our goal is to estimate x_1 and x_2 as well as parameter α using the proposed nonlinear censored filter. The estimation of model parameters with Kalman filtering has received considerable attention. A popular approach is the so-called dual estimation method [33, 34, 35, 20] which treats the model parameters q as auxiliary state variables that evolve slowly over time. In this article we assume persistent dynamics, namely $\dot{q} = 0$. Using this approach, we assume $\dot{\alpha} = 0$ and form an augmented state vector consisting of the original state variables x_1 and x_2 and now α, thus allowing for simultaneous state and parameter estimation.

Fig. 1 shows the estimation results when $\alpha = 1$ and the observations of x_1 are corrupted by 30% Gaussian observational noise. Black circles indicate the noisy

observations, dotted black lines denote the true trajectory of the variables and parameters and solid grey curves reflect the filter estimate. In the estimation results for parameter α, we also include the filter estimated 95% confidence interval (dashed grey curves). After an initial transient period, the filter is able to estimate the system variables and parameter with great accuracy. Of particular importance, we notice the fidelity of the filter reconstruction of the variables during the periods of censored data.

As previously mentioned, one of the main advantages of using the Kalman filter for estimation is that it is a sequential estimator. While this means that new observations can be processed online without re-analzying the entire dataset, the more important implication is that it allows for the tracking of parameters whose values may drift over time. To simulate this scenario, we considered the estimation problem in the above oscillator system when α changes over time. Namely, its value changes from 1 to 0.5 after 15 units of time. Again, we work under the assumption that only observations of x_1 affected by 30% observational noise are available and also $\dot{\alpha} = 0$. Fig. 2 shows the resulting estimation in this nonstationary case. Once again after the initial transient period of the filter we see convergence of α to its correct value and accurate estimation of the state variables. As α drifts, the filter loses track of the x_1 and x_2 variables but is able to recover after a sufficient amount of data has been observed. Furthermore, the filter is able to accurately track the drift in α, despite the presence of censored data.

4. Estimation in HCV System

Given the success of the nonlinear censored filter in the oscillator system above, we now consider a significantly more difficult example of state and parameter estimation for analyzing HCV patient data. In this example, an HCV-infected liver undergoes antiviral treatment with interferon-α (IFN) and ribavirin. The typical measurement in this clinical setting is the patient's viral load. Unfortunately viral load is only detectable above a threshold of about 50

copies/mL [36]. This means that any measurements below this level are censored (i.e. our censored interval in this case would be $[-\infty, 50]$). Several HCV models have developed, and in particular we consider one by Snoeck, et al. [36]. This system is described by the following equations

$$\begin{aligned}
\frac{dT}{dt} &= s + rT\left(1 - \frac{T+I}{T_{max}}\right) - dT - \beta V_I T \\
\frac{dI}{dt} &= \beta V_I T + rI\left(1 - \frac{T+I}{T_{max}}\right) - \delta I \\
\frac{dV_I}{dt} &= (1-\bar{\rho})(1-\bar{\epsilon})pI - cV_I \\
\frac{dV_{NI}}{dt} &= \bar{\rho}(1-\bar{\epsilon})pI - cV_{NI},
\end{aligned} \tag{6}$$

where T and I denote concentrations of healthy and infected hepatocytes, and V_I and V_{NI} denote concentrations of infectious and noninfectious virions, respectively. Of note, for parameters $\bar{\rho}$ and $\bar{\epsilon}$ we assume exponentially decaying dynamics

$$\begin{aligned}
\bar{\rho} &= \rho e^{-k(t-t_{end})_+} \\
\bar{\epsilon} &= \epsilon e^{-k(t-t_{end})_+},
\end{aligned}$$

where t_{end} indicates the end of treatment and

$$(a)_+ = \begin{cases} a & \text{if } a \geq 0 \\ 0 & \text{if } a < 0. \end{cases}$$

With respect to (6), the viral load data maps to the quantity $y = V_I + V_{NI}$. The state variables and parameters of (6) can range over many orders of magnitude, making accurate estimation difficult. To aid in this process, we apply transformations to all estimated components. In particular, we apply a $\log_{10}$ transformation to the state vector x to compute $\tilde{x}$ and use the relationship

$$\frac{d\tilde{x}_j}{dt} = \frac{1}{\ln(10)x_j}\frac{dx_j}{dt}$$

for all j, where $x_j = 10^{\tilde{x}_j}$. We also scale the parameters as $\tilde{q}_j = \log_{10} q_j$ for all q_j except the efficacy values ϵ and ρ. These two parameters must be constrained

to the interval $[0, 1]$, so we instead use

$$\tilde{q}_j = \tan\left(\pi q_j - \pi/2\right).$$

With this log transformation of the model, we assume our filter observation function h to be

$$h = \log_{10}\left(V_I + V_{NI}\right)$$

The numerous parameters in the model, combined with the limited (and often censored) observability of the physical system, presents a difficult estimation problem. A thorough consideration of this HCV model, including parameter senstivity analysis and estimation, was considered in [26]. There, the authors used expectation maximization (EM) to estimate the model states and identifiable parameters for different HCV datasets. Here, we assume the results of [26] to be our "ground truth" and attempt to show that the nonlinear censored filter is able to converge to similar estimates. Again, we emphasize that the censored filter should not give us better or more accurate results than EM, but rather an alternative approach that allows for the sequential estimation of state variables and parameters.

For a full description of the model variables and parameters, see [36, 26]. Here, we restrict ourselves to the analysis of data from a patient in relapse as found in [36]. Fig. 3 shows the log-scaled viral load measurements (black circles) from said patient. We notice immediately that there is a clear lower limit of detection, resulting in censored observability of the system. Following the analysis of [26], we fix the parameter values detailed in Table 1. Using the censored filter, we estimate the transformed state variables $\tilde{T}$, $\tilde{I}$, $\tilde{V}_I$ and $\tilde{V}_{NI}$ and parameters $\tilde{\delta}$, $\tilde{c}$ and $\tilde{\epsilon}$ which correspond to the infected cell death rate, virion elimination rate and peginterferon efficacy respectively. Like in our previous example, we assume persistent dynamics for the parameters of interest which allows us to form an augmented state vector and implement the dual-estimation scheme.

Fig. 4 shows the results of the HCV parameter estimation for analyzed pa-

tient data. Dotted black lines denotes the converged estimate as found using EM in [26] and grey curves correspond to the censored filter estimate. Dotted grey lines indicate the filter estimated 95% confidence region of the estimates. After a sufficient amount of data, the censored filter is able to converge to parameters estimates comparable to that of EM. Fixing the estimated parameters to their converged estimates, we re-run the filter to obtain an accurate estimation of the state variables. Fig. 5 shows the final log-scaled viral load estimate. We obtain a good fit of the data and furthermore we are able to get a reasonable estimation of the system state during the censored data regions.

5. Estimation in HIV System

We now conisder a more sophisticated example of studying in-host HIV dynamics. The patient data analyzed here comes from a clinical study at Massachusetts General Hospital between 1996 and 2004. This data, originally examined in [37], consists of two measured quantities: $CD4^+$ T-lymphocyte count (cells/μL) and viral load (copies/mL). Measurement of viral load is once again subject to the detection limits of the assay. In this study both a standard assay, with a detection limit of 400 copies/mL and above, and a high sensitivity assay, with a detection limit of 50 copies/mL and above, were used. Any measurements below the detection limits of the respective assays were effectively censored.

A complex model of in-host HIV dynamics developed in [38] is described by

the following system of equations

$$
\begin{aligned}
\dot{T}_1 &= \lambda_1 - d_1 T_1 - (1-\epsilon_1(t)) d_1 V_I T_1 \\
\dot{T}_2 &= \lambda_2 - d_2 T_2 - (1-f\epsilon_1(t)) k_2 V_I T_2 \\
\dot{T}_1^* &= (1-\epsilon_1(t)) k_1 V_I T_1 - \delta T_1^* - m_1 T_1^* E \\
\dot{T}_2^* &= (1-f\epsilon_1(t)) k_2 V_I T_2 - \delta T_2^* - m_2 T_2^* E \\
\dot{V}_I &= (1-\epsilon_2(t)) N_T \delta (T_1^* + T_2^*) \\
&- (c + (1-\epsilon_1(t))\rho_1 k_1 T_1 + (1-f\epsilon_1(t))\rho_2 k_2 T_2) V_I \\
\dot{V}_{NI} &= \epsilon_2(t) N_T \delta (T_1^* + T_2^*) - c V_{NI} \\
\dot{E} &= \lambda_E + b_E \frac{T_1^* + T_2^*}{T_1^* + T_2^{8} + K_b} E \\
&- d_E \frac{T_1^* + T_2^*}{T_1^* + T_2^* + K_d} E - \delta_E E.
\end{aligned}
\tag{7}
$$

The model state variables consist of T_1 (uninfected type 1 target cells, e.g. $CD4^+$ T-cells), T_2 (uninfected type 2 target cells, e.g. magrophages), T_1^* (infected type 1 target cells), T_2^* (infected type 2 target cells), V_I (infectious free virus), V_{NI} (non-infectious free virus) and E (cytotoxic T-lymphocytes, e.g. CD8 cells). The units for the model variables are in μL. Treatment is modeled through $\epsilon_1(t) = \epsilon_1 u(t)$ and $\epsilon_2(t) = \epsilon_2 u(t)$ where $0 \leq u(t) \leq 1$.

An example of the data collected from a patient in the study is shown in Fig. 6. We notice that the measurement of $CD4^+$ and viral load often occur at different intervals. Additionally, we observe a clear lower limit for viral load detection. With regards to (7), the collected $CD4^+$ data maps to quantity $y_1 = T_1 + T_1^*$ and the collected viral load data maps to $y_2 = V_I + V_{NI}$. For a detailed description of (7) and the estimation analysis done for the data acquired in the clinical study, including patient-specific identifiability analysis, see [27]. Once again, our goal is merely to show that the censored filter derived here is able to reconstruct similar state variable and parameter estimates as those found in [27] which used the established EM method. We restrict our investigation to the patient data shown in Fig. 6.

Similarly to the HCV model discussed in the previous section, the HIV model

variables and parameters can vary on drastically different orders of magnitude. As such, we once again introduce a $\log_{10}$ transformation for the model variables and parameters to allow for a more robust estimation procedure. The observation function for the filter consists of quantities $h_1 = \log_{10}(T_1 + T_1^*)$ and $h_2 = \log_{10}(V_I + V_{NI})$. However as mentioned earlier, the data are collected at different intervals meaning that the filter's observation function changes with respect to the data available at each assimilation time point.

Given these observations, our goal is to estimate log-scaled variables $\tilde{T}_1$, $\tilde{T}_2$, $\tilde{T}_1^*$, $\tilde{T}_2^*$, $\tilde{V}_I$, $\tilde{V}_{NI}$, $\tilde{E}$ and log-scaled parameters $\tilde{k}_1$ and $\tilde{k}_2$ which correspond to the population 1 and population 2 infection rates respectively. As in our previous examples, we assume persistent dynamics for the parameters thereby allowing us to implement dual estimation. Parameters that were not estimated were fixed to the values in Table 2 as detailed in [27].

Fig. 7 shows the results of the filter estimation for transformed parameters $\tilde{k}_1$ and $\tilde{k}_2$. After a sufficient amount of data, the filter estimates (solid grey curves) converge to the desired parameter values (dotted black curve) that were obtained using EM. Additionally, the estimated 95% confidence region (dashed grey curves) shrinks as convergence occurs reinforcing the optimality of the parameter estimate. We once again fix the estimated parameters to their convergent values and re-run the filter to obtain accurate estimates of the state variables. The resulting filter estimates (solid grey curves) are shown in Fig. 8. We obtain a good fit of both data and additionally are able to get a reasonable reconstruction of the viral load during the regions of censored data, once again showing the capabilities of the censored filter.

6. Conclusion

The presence of censored data further complicates the state and parameter estimation process. Incorrectly accounting for these observations can lead to inaccurate estimates resulting in incorrect model inference. Here we derived a modified version of the extended Kalman filter which accounts for the cen-

sored observations in the form of an auxiliary set of filter update equations. We examined the performance of this novel filter in an oscillator system where measurements were noisy and censored. We demonstrated its ability to accurately reconstruct state variables and track stationary and drifting parameter values despite the limitations imposed by the censored data. Motivated by this success, we implemented the filter to analyze complex censored data from an HCV and HIV clinical study. The proposed filter was able to obtain comparable estimates for the parameters and state variables as those calculated in the literature using expectation maximization.

The success of the nonlinear censored filter opens up many exciting possibilities. The sequential nature of the algorithm allows for the online estimation of states and parameters and more importantly the tracking of parameters within a patient's dataset that may change over time. Being able to track any potential parameter drift would allow for much more accurate model-based prescription of treatment. Further work should examine the implementation of other nonlinear filters, such as the unscented and ensemble Kalman filter, in place of the EKF which can be costly due to the required system linearization.

7. Acknowledgments

This research was partially supported by grants No. RTG/DMS-1246991 and No. DMS-1022688 from the National Science Foundation.

8. Appendix A. Conditional Moment Calculations

Suppose x and z are m- and n-dimensional, jointly Gaussian random vectors. Additionally, let $\mathcal{Z}$ be an n-dimensional rectangle in $\mathbb{R}^n$. Then the conditional mean of x given $z \in \mathcal{Z}$ is

$$\begin{aligned}\mathbf{E}\left[x|z \in \mathcal{Z}\right] =& \mathbf{E}\left[\mathbf{E}\left[x|z\right]|z \in \mathcal{Z}\right] \\ =& \mathbf{E}\left[\mu_x + K(z-\mu_z)|z \in \mathcal{Z}\right] \\ =& \mu_x + K(\mu_{z|z\in\mathcal{Z}} - \mu_z),\end{aligned}$$

where $K = P_{xz}P_z^{-1}$ and $\mu_{z|z\in\mathcal{Z}}$ is the conditional mean of z given $z \in \mathcal{Z}$ [1]. The derivation of the conditional covariance is more lengthy, but the result has the simple form

$$P_{x|z\in\mathcal{Z}} = P_x - K(P_z - P_{z|z\in\mathcal{Z}})K^T,$$

where $P_{z|z\in\mathcal{Z}}$ is the covariance of z conditional on $z \in \mathcal{Z}$ [1].

9. Appendix B. Covariance Prediction

Consider the covariance $D(t)$ between C, the vector of censored observations, and the current state vector $x(t)$. This is

$$\begin{aligned} D(t) =& \mathbf{E}\left[(C - \hat{C})(x(t) - \hat{x}(t))\right] \\ =& \mathbf{E}\left[Cx(t)\right] - \hat{C}\hat{x}(t), \end{aligned}$$

where, keeping with our censored data Kalman filter, $\hat{C}$ and $\hat{x}(t)$ are expectations given the uncensored data. We are interested in how $D(t)$ evolves during a time interval when there are no new measurements. Omitting the explicit time-dependence for x and D, we have

$$\begin{aligned} \frac{d}{dt}D =& \frac{d}{dt}\left(\mathbf{E}\left[Cx^T\right] - \hat{C}\hat{x}^T\right) \\ =& \frac{d}{dt}\mathbf{E}\left[Cx^T\right] - \hat{C}\frac{d}{dt}\hat{x}^T. \end{aligned}$$

The first term can be simplified as

$$\begin{aligned} \frac{d}{dt}\mathbf{E}\left[Cx^T\right] =& \mathbf{E}\left[C\frac{d}{dt}x^T\right] \\ =& \mathbf{E}\left[C\left(f(x) + g(t)w(t)\right)^T\right] \\ =& \mathbf{E}\left[Cf(x)^T\right] \\ \approx& \mathbf{E}\left[C\left(f(\hat{x}) + \nabla f(\hat{x})(x - \hat{x})\right)^T\right] \\ =& \hat{C}f(\hat{x})^T + \mathbf{E}\left[Cx^T\right]\nabla f(\hat{x})^T - \hat{C}\hat{x}^T\nabla f(\hat{x})^T, \end{aligned}$$

where we have used the fact that $w(t)$ is independent of C and has expectation 0. Subtracting off $\hat{C}\frac{d}{dt}\hat{x}^T$ with the substitution $\frac{d}{dt}\hat{x}^T \approx f(\hat{x})^T$ we have

$$\begin{aligned}\frac{d}{dt}D \approx & \mathbf{E}\left[Cx^T\right]\nabla f(\hat{x})^T - \hat{C}\hat{x}^T\nabla f(\hat{x})^T \\ = & D\nabla f(\hat{x})^T.\end{aligned}$$

References

References

[1] B. Ibarz-Gabardos, P. Zufiria, A kalman filter with censored data, in: Proceedings IEEE International Workshop on Intelligent Signal Processing, IEEE, 2005, pp. 74–79.

[2] R. Kalman, A new approach to linear filtering and prediction problems, J. Basic Eng. 82 (1960) 35–45.

[3] E. Kalnay, Atmospheric modeling, data assimilation, and predictability, Cambridge Univ. Press, 2003.

[4] G. Evensen, Data assimilation: The Ensemble Kalman Filter, Springer: Heidelberg, 2009.

[5] F. Rabier, Overview of global data assimilation developments in numerical weather-prediction centres, Quarterly Journal of the Royal Meteorological Society 131 (613) (2005) 3215–3233.

[6] B. Hunt, E. Kalnay, E. Kostelich, Four-dimensional ensemble Kalman filtering, Tellus A 56 (2004) 273–277.

[7] J. A. Cummings, Operational multivariate ocean data assimilation, Quarterly Journal of the Royal Meteorological Society 131 (613) (2005) 3583–3604. doi:10.1256/qj.05.105.

[8] K. Yoshida, J. Yamaguchi, Y. Kaneda, Regeneration of small eddies by data assimilation in turbulence, Phys. Rev. Lett. 94 (2005) 014501.

[9] K. Law, A. Stuart, Evaluating data stimulation algorithms, Mon. Wea. Rev. 140 (2012) 3757–3782.

[10] S. Schiff, Neural control engineering, MIT Press, 2012.

[11] T. Berry, T. Sauer, Adaptive ensemble kalman filtering of nonlinear systems, Tellus A 65 (2013) 20331.

[12] G. Ullah, S. Schiff, Tracking and control of neuronal Hodgkin-Huxley dynamics, Phys. Rev. E 79 (2009) 040901.

[13] T. Sauer, S. Schiff, Data assimilation for heterogeneous networks: The consensus set, Phys. Rev. E 79 (2009) 051909.

[14] G. Ullah, S. Schiff, Assimilating seizure dynamics, PLoS Computational Biology 6 (2010) e1000776.

[15] F. Hamilton, T. Berry, N. Peixoto, T. Sauer, Real-time tracking of neuronal network structure using data assimilation, Physical Review E 88 (2013) 052715.

[16] F. Hamilton, J. Cressman, N. Peixoto, T. Sauer, Reconstructing neural dynamics using data assimilation with multiple models, Europhysics Letters 107 (2014) 68005.

[17] H. Dahari, A. Lo, R. Ribeiro, A. Perelson, Modeling hepatitis c virus dynamics: Liver regeneration and critical drug efficacy, Journal of theoretical biology 247 (2) (2007) 371–381.

[18] A. Neumann, N. Lam, H. Dahari, D. Gretch, T. Wiley, T. Layden, A. Perelson, Hepatitis c viral dynamics in vivo and the antiviral efficacy of interferon-α therapy, Science 282 (5386) (1998) 103–107.

[19] A. Perelson, A. Neumann, M. Markowitz, J. Leonard, D. Ho, et al., Hiv-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time, Science 271 (5255) (1996) 1582–6.

[20] H. Voss, J. Timmer, J. Kurths, Nonlinear dynamical system identification from uncertain and indirect measurements, Int. J. Bif. Chaos 14 (2002) 1905–1924.

[21] A. Sitz, U. Schwarz, J. Kurths, H. Voss, Estimation of parameters and unobserved components for nonlinear systems from noisy time series., Physical Review E 66 (2002) 16210.

[22] B. Matzuka, Nonlinear filtering methodologies for parameter estimation and uncertainty quantification in noisy, complex, biological systems, Ph.D. thesis.

[23] B. Matzuka, M. Aoi, A. Attarian, H. Tran, Nonlinear filtering methodologies for parameter estimation, Tech. rep., North Carolina State University CRSC (2012).

[24] B. Allik, C. Miller, M. Piovoso, R. Zurakowski, Nonlinear estimators for censored data: a comparison of the ekf, the ukf and the tobit kalman filter, in: Proceedings American Control Conference, IEEE, 2015, pp. 5146–5151.

[25] B. Allik, C. Miller, M. Piovoso, R. Zurakwoski, The tobit kalman filter: an estimator for censored measurements, IEEE Transactions on Control Systems Technology 24 (2015) 365–370.

[26] J. Arthur, H. Tran, P. Aston, Feasibility of parameter estimation in hepatitis c viral dynamics models, Journal of Inverse and Ill-Posed Problems In Press.

[27] A. Attarian, Patient specific subset selection, estimation and validation of an hiv-1 model with censored observations under an optimal treatment schedule, Ph.D. thesis (2012).

[28] M. Geoffrey, K. Thriyambakam, The EM algorithm and extensions, John Wiley and Sons, 2008.

[29] D. Simon, Optimal State Estimation: Kalman, H_∞, and Nonlinear Approaches, John Wiley and Sons, 2006.

[30] K. Law, A. Stuart, K. Zygalakis, Data Assimilation: A Mathematical Introduction, Springer, 2015.

[31] J. David, H. Tran, H. Banks, Hiv model analysis and estimation implementation under optimal control based treatment strategies, International Journal of Pure and Applied Mathematics 57 (3) (2009) 357–392.

[32] S. Wilhelm, M. B. G, tmvtnorm: Truncated Multivariate Normal and Student t Distribution, r package version 1.4-8 (2013).
URL `http://CRAN.R-project.org/package=tmvtnorm`

[33] H. Cox, On the estimation of state variables and parameters for noisy dynamic systems, IEEE Trans. Automatic Control 9 (1964) 5–12.

[34] R. Kopp, R. Orford, Linear regression applied to system identification for adaptive control systems, AIAA Journal 1 (1963) 2300–2306.

[35] E. Wan, A. Nelson, Dual extended Kalman Filter methods, Kalman Filtering and Neural Networks, Ed. S. Haykin (2001) 123–173.

[36] E. Snoeck, P. Chanu, M. Lavielle, P. Jacqmin, E. Jonsson, K. Jorga, T. Goggin, J. Grippo, N. Jumbe, N. Frey, A comprehensive hepatitis c viral kinetic model explaining cure, Clinical Pharmacology & Therapeutics 87 (6) (2010) 706–713.

[37] B. Adams, Non-parametric parameter estimation and clinical data fitting with a model of hiv infection, Ph.D. thesis (2005).

[38] B. Adams, H. T. Banks, M. Davidian, H. dae Kwon, H. Tran, S. Wayne, E. Rosenberg, Hiv dynamics: modeling, data analysis and optimal treatment protocols, Journal of Computational and Applied Mathematics 184 (2005) 10–49.

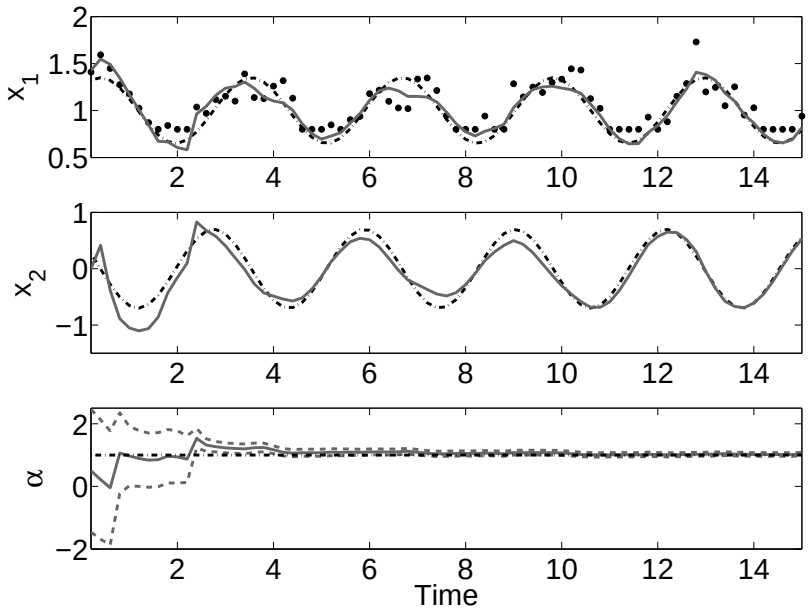

Figure 1: State and parameter estimation in oscillator system when α is fixed over time. Observations (black circles) of the x_1 variable are perturbed by 30% observational noise and censored below a value of 0.8. Dotted black lines denote the true variable/parameter trajectory and solid grey curves the filter estimates. Dashed grey curve denote the filter estimated 95% confidence region. Despite the presence of censored data, the filter is able to accurately estimate both state variables as well as the unknown parameter. In particular, we note the fidelity of the reconstruction during censored regions of the data.

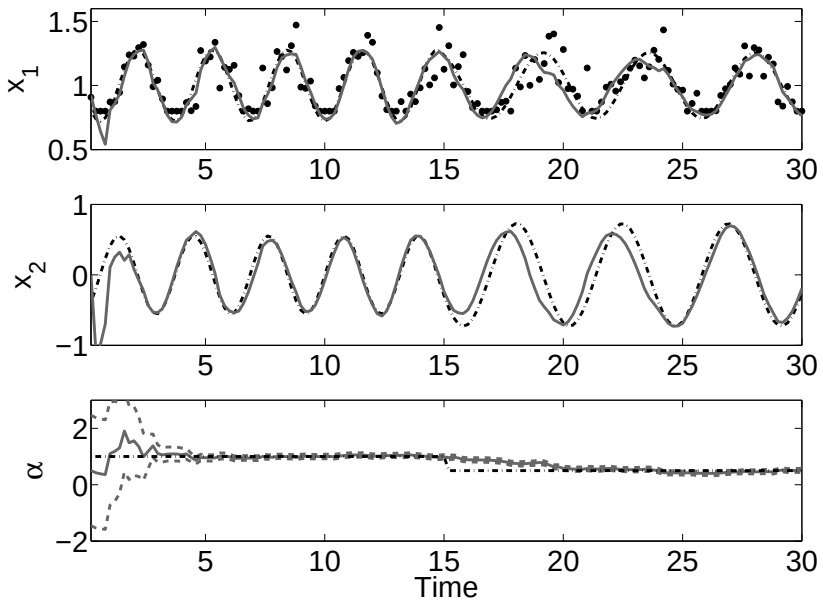

Figure 2: State and parameter estimation in simple system when α drifts over time. Observations (black circles) of x_1 are perturbed by 30% Gaussian observational noise. Dotted black lines indicate the true variable/parameter trajectory and solid grey curves the filter estimates. Once again we include the filter estimated 95% confidence region (dashed grey curve) for α. In this more complicated example where a system nonstationarity is present, the filter is still able to accurately track the drift in α and reconstruct the state variables even when there are censored data.

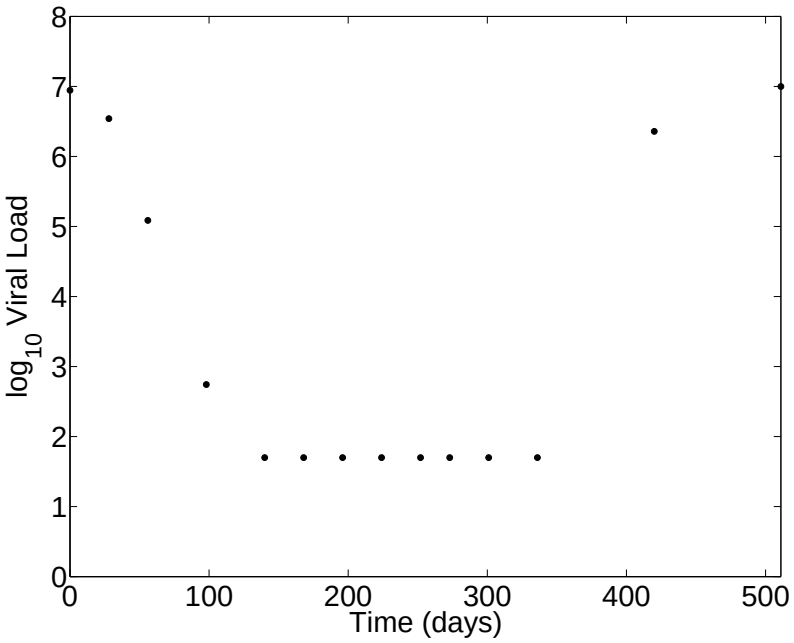

Figure 3: Log-scaled viral load data (black circles) collected from a patient in relapse. We observe a clear detection limit in the measurement of viral load, leading to a censored estimation problem.

Table 1: Fixed Parameter Values for HCV Patient Data

Parameter	Description	Value
β	Infection rate	8.7×10^{-9}
p	Virion production rate	25.1
r	Cell proliferation rate	5.620×10^{-3}
ρ	Ribavirin efficacy	0.5
k	Efficacy decay rate	0.0238
s	Cell production rate	6.17×10^{4}
T_{max}	Total number of cells per mL	1.85×10^{7}
d	Cell death rate	0.003

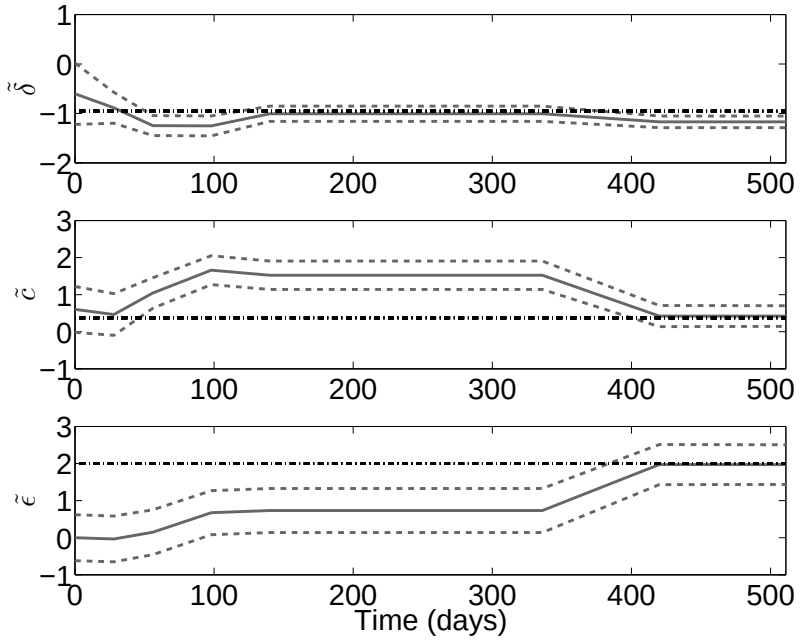

Figure 4: Results for the estimation of the transformed HCV model parameters $\tilde{\delta}$, $\tilde{c}$ and $\tilde{\epsilon}$ in relapse data set. Dotted black lines denote "ground truth" transformed parameter values found in [26] and grey solid lines indicate the censored filter estimate. Dotted grey lines indicate the filter estimated 95% confidence region of the estimate. After an initial transient period, the filter is able to converge to the estimates obtained using EM.

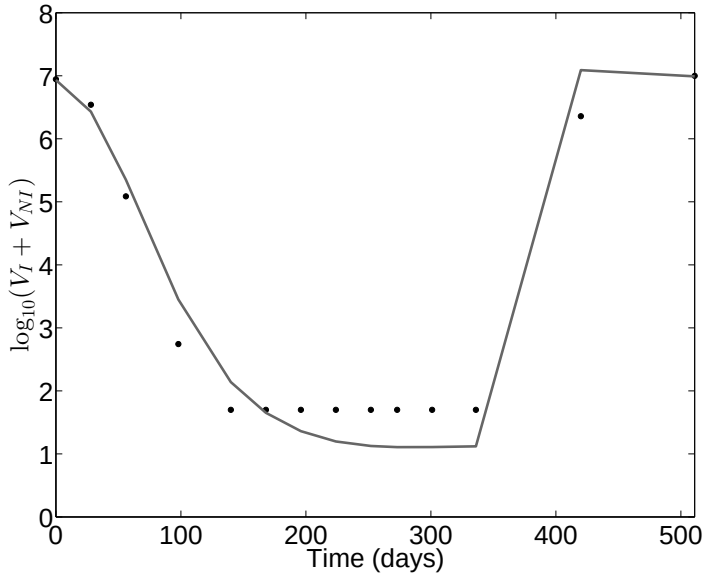

Figure 5: Estimation of log-scaled viral load when estimated parameters are fixed to their converged values. Observations (black circles) and filter estimate (grey curve) shown. We obtain a reasonable fit for the data and furthermore estimate a smooth trajectory for the viral load during the censored region of the data.

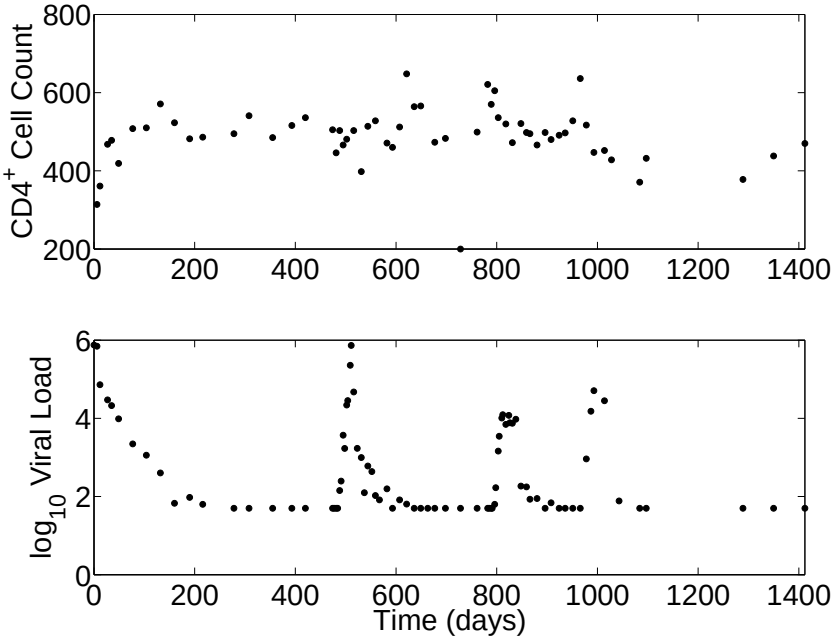

Figure 6: Example of patient data (black circles), $CD4^+$ cell count and log-scaled viral load, from clinical study shown. Of note, the measurement of $CD4^+$ and viral load often occur at different time intervals. Furthermore, there is a clear detection limit for measurement of viral load resulting in a censored estimation problem.

Table 2: Fixed Parameter Values for HIV Patient Data

Parameter	Description	Value
λ_1	Target cell type 1 source rate	4.4111
λ_2	Target cell type 2 source rate	0.0342
d_1	Target cell type 1 death rate	9.91029×10^{-3}
d_2	Target cell type 2 death rate	2.6601×10^{-3}
m_1	Population 1 immune-induced clearance rate	2.8674×10^{-6}
m_2	Population 2 immune-induced clearance rate	2.9136×10^{-6}
ρ_1	Virions infecting type 1 cell	0.99052
ρ_2	Virions infecting type 2 cell	0.99622
δ	Infected cell death rate	0.0952
c	Virus death rate	11.4004
f	Treatment efficacy reduction in population 2	0.0980
N_T	Virions produced per infected cell	102.5980
λ_E	Immune effector source rate	9.4159×10^{-4}
δ_E	Immune effector death rate	0.1201
b_E	Immune effector max birth rate	0.0826
d_E	Immune effector max death rate	0.0939
K_b	Saturation constant for immune effector birth	0.1082
K_d	Saturation constante for immune effector death	0.1009
ϵ_1	Reverse transcriptase inhibitor efficacy	0.5140
ϵ_2	Protease inhibitor efficacy	0.5770

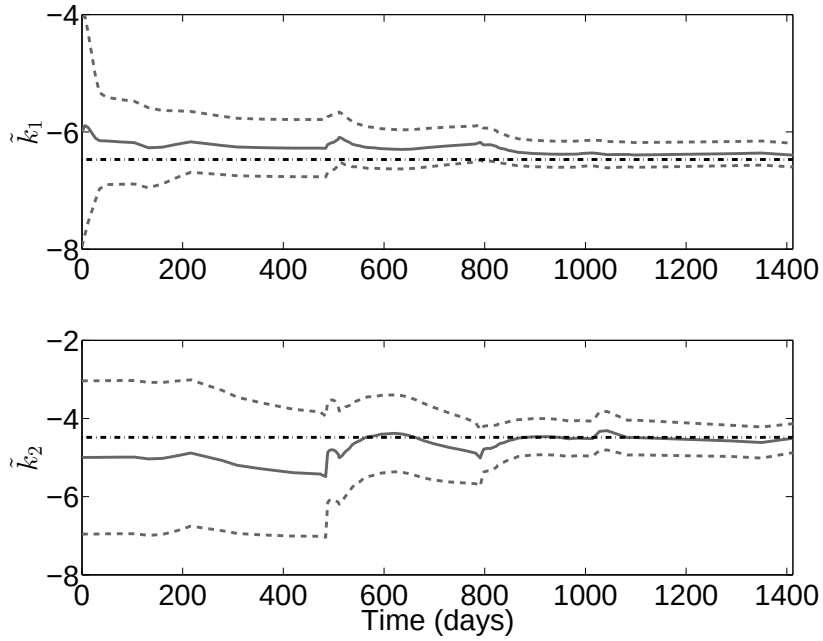

Figure 7: Estimated log-transformed parameters for HIV patient dataset. Filter estimated parameter values (grey curve) compare favorably with the values estimated by EM (dotted black line). Filter estimated 95% confidence interval also shown (dotted grey curves). As the estimates converge to the correct value, the confidence interval shrinks showing reliability of estimates.

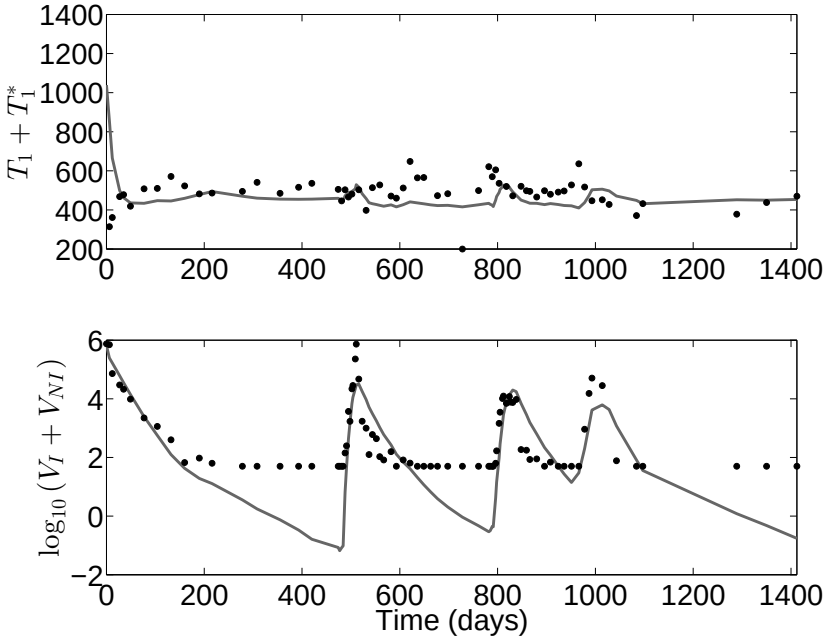

Figure 8: Estimation of log-scaled viral load and CD4$^+$ count when estimated parameters are fixed to their converged values. Observations (black circles) and filter estimate (grey curve) shown. We obtain good fits for the data and in particular we obtain a good reconstruction of the viral load during the censored regions of the data.